

10/758,485

IBM/286

In the Claims

1. (Presently amended) A method of performing a query in a relational database system by operating upon a plurality of relations each comprising a plurality of tuples formed over a plurality of attributes, comprising:

evaluating join predicates in said query to determine whether a join involving a first relation and a second relation will be reductive of said first relation, and if so, identifying a join involving said first and second relations that will be reductive of said first relation, and performing said query by the prior application of a look-ahead predicate based upon the second relation in the join.

2. (original) The method of claim 1 further comprising determining whether said second relation involved in the join is subject to a selection criterion, and evaluating whether that selection criterion effects a join reduction.

3. (original) The method of claim 2 wherein an amount of join reduction effected by a selection criterion is determined by identifying whether the number of rows in the join result will be smaller than the number of rows in the first relation.

10/758,485

IBM/286

4. (original) The method of claim 2 wherein, upon identifying a join reduction involving a first and a second relation, and a selection criterion on the second relation, the potential benefit of that join reduction is assessed.

5. (original) The method of claim 4 further comprising evaluating the computational expense of generating a look-ahead predicate comprising the tuples of the second relation matching the selection criterion, and comparing said expense to computational savings that result from the join reduction.

6. (original) The method of claim 5 further comprising, upon identifying a beneficial look-ahead predicate, processing the query by forming and utilizing the look-ahead predicate as a selection criterion on the second relations.

7. (original) The method of claim 5 further comprising identifying the most beneficial look-ahead predicate among all potential joins of relations in said query, through iterative analysis of all possible joins.

8. (original) The method of claim 7 further comprising iteratively analyzing all possible joins of the remaining relations and the look-ahead predicate to locate further beneficial look-ahead predicates.

10/758,485

IBM/286

9. (Presently amended) Apparatus for performing a query in a relational database system by operating upon a plurality of relations each comprising a plurality of tuples formed over a plurality of attributes, comprising:

a data storage device storing said relations, and

a processor evaluating join predicates in said query to determine whether a join involving a first relation and a second relation will be reductive of said first relation, and if so, identifying a join involving said first and second relations that will be reductive of said first relation, and performing said query by the prior application of a look-ahead predicate based upon the second relation in the join.

10. (original) The apparatus of claim 9 wherein said processor determines whether a relation involved in the join is subject to a selection criterion, and evaluates whether that selection criterion effects a join reduction.

11. (original) The apparatus of claim 10 wherein an amount of join reduction effected by a selection criterion is determined by identifying whether the number of rows in the join result will be smaller than the number of rows in the first relation.

12. (original) The apparatus of claim 10 wherein, upon identifying a join reduction involving a first and a second relation, and a selection criterion on the second relation, said processor assesses the potential benefit of that join reduction.

10/758,485

IBM/286

13. (original) The apparatus of claim 12 wherein said processor evaluates the computational expense of generating a look-ahead predicate comprising the tuples of the second relation matching the selection criterion, and comparing said expense to computational savings that result from the join reduction.

14. (original) The apparatus of claim 13 wherein, upon identifying a beneficial look-ahead predicate, said processor processes the query by forming and utilizing the look-ahead predicate as a selection criterion on the second relations.

15. (original) The apparatus of claim 13 wherein said processor identifies the most beneficial look-ahead predicate among all potential joins of relations in said query, through iterative analysis of all possible joins.

16. (original) The apparatus of claim 15 wherein said processor iteratively analyzes all possible joins of the remaining relations and the look-ahead predicate to locate further beneficial look-ahead predicates.

17. (Presently Amended) A program product comprising:

a relational database comprising one or more relations, each relation comprising one or more tuples on one or more attributes, and

10/758.485

IBM/286

relational database system adapted to perform a query on said relational database by evaluating join predicates in said query to determine whether a join involving a first relation and a second relation will be reductive of said first relation, and if so, identify a join involving said first and second relations that will be reductive of said first relation, and perform said query by the prior application of a look-ahead predicate based upon the second relation in the join, and signal bearing media bearing the relational database and the relational database system.

18. (original) The program product of claim 17 wherein the signal bearing media comprises transmission media.

19. (original) The program product of claim 17 wherein the signal bearing media comprises recordable media.